4) TYPER - For the AIM workshop held at Stanford last summer, we developed a program to assist with on-line typescript manipulation. The goal of the workshop (see Section I.E.1 on page 79) was to present a deeper insight into the workings of developing AI programs by interactively tracing sessions using them. In order to assure a reasonably organized presentation that could be prepared beforehand, we developed TYPER to allow a presenter to display a typescript of a typical session. TYPER provides facilities to randomly move between parts of the typescript, to display a table of contents, and to manage a hierarchical presentation of various parts of the session. At the highest level, only an overview of program operation need be given. By interactive commands, successive layers of detail can be flashed on the display screen as the discussion proceeds or as questions arise.

We have also implemented extensions and maintenance updates to many other existing programs including, for example, EDIR (a directory editing program), DUMPER (the file system backup program), BSYS (the user file archiving program), PA-1050 (the TOPS-10 compatibility package), BBD (the bulletin board program), and PUB (a text formatting program).

I.A.3.5 Documentation and Education

We have spent considerable effort to develop, maintain, and facilitate access to our documentation so as to accurately reflect available software. The HELP and Bulletin Board subsystems have been important in this effort. As subsystems are updated, we generally publish a bulletin or small document describing the changes. As more and more changes occur, it becomes harder and harder for users to track down all of the change pointers. Within manpower limits, we are in a continuous process of reviewing the existing documentation system for compatibility with the programs now on line and to integrate changes into the main documents. This will also be done with a view toward developing better tools for maintaining up-to-date documentation.

I.A.3.6 Software Compatibility and Sharing

At SUMEX-AIM we firmly believe in importing rather than reinventing software where possible. As noted above, a number of the packages we have brought up are from outside groups. Many avenues exist for sharing between the system staff, various user projects, other facilities, and vendors. The advent of fast and convenient communication facilities coupling communities of computer facilities has made possible effective intergroup cooperation and decentralized maintenance of software packages. The TENEX sites on the ARPANET have been a good model for this kind of exchange based on a functional division of labor and expertise. The other major advantage is that as a by-product of the constant communication about particular software, personal connections between staff members of the various sites develop. These connections serve to pass general information about software tools and to encourage the exchange of ideas among the sites.

Certain common problems are now regularly discussed on a multi-site level. We continue to draw significant amounts of system software from other ARPANET sites, reciprocating with our own local developments. Interactions have included mutual backup support, experience with various hardware configurations, experience with new types of computers and operating systems, designs for local networks, operating system enhancements, utility or language software, and user project collaborations. We have been able to import many new pieces of software and improvements to existing ones in this way. Examples of imported software include the message manipulation program MSG, TENEX SAIL, PASCAL, TENEX SOS, INTERLISP, the RECORD program, ARPANET host tables, and many others. Reciprocally, we have exported our contributions such as the crash analysis program, drum page migration system, KI-10 page table efficiency improvements, GTJFN enhancements, PUB macro files, the bulletin board system, MAINSAIL, SPELL, SNDMSG enhancements, our BATCH monitor, and improved SA-10 software.

There are also several important examples of joint development efforts such as the internet mailer program (XMAILR). Because this program incorporates facilities for routing mail through many networks, it is important that the various sections of the program dealing with these specialized protocols be developed by the groups with expertise in the appropriate technology. Network connections have made a joint effort possible involving MIT, Stanford SCORE, and SUMEX.

We spent considerable effort developing a preliminary version of a TENEX/TOPS-20 compatibility package. The issue here is that as DEC develops TOPS-20, even though it is TENEX-like, it is not TENEX compatible and vice versa. Thus, a hope was to write a program that would resolve these compatibilities automatically rather than to force special adaptations for the two operating systems. The kinds of incompatibilities that exist include PDP-20 machine instructions that do not exist on earlier machines, new JSYS calls, incompatible changes to old JSYS calls, different syntax and facilities for device/file names, and different handling of error returns (types of return and error codes). It has proven unworkable to effectively handle all of these problems at the user level. Monitor changes are required to implement the widely used error return features (ERJMP/ERCAL) and make handling of other incompatibilities easier. We have not accomplished a complete compatibility package in any sense but have implemented requisite monitor changes and have developed several user packages that help emulate TOPS-20 JSYS calls for programs running on TENEX. We do not foresee being able to completely and effectively solve these problems within the expected lifetime of existing TENEX machines.

Finally, we have also assisted groups that have interacted with SUMEX user projects get access to software available in our community. For example, Prof. Dreiding's group in Switzerland became interested in some of the system software available here after attending the DENDRAL CONGEN workshops (see Section II.A.1.3 on page 103). We have provided him with the non-licensed programs requested. We have also provided software to Professor Bodmer's group at the Imperial Cancer Research Group in England in collaboration with the MOLGEN project (see Section II.A.1.5 on page 136).

I.A.3.7 Core Research

Over the past year we have supported several core research activities aimed at developing information resources, basic AI research, and tools of general interest to the SUMEX-AIM community. Principal areas of current effort include:

1) The AI Handbook which is a compendium of knowledge about the field of Artificial Intelligence being compiled by Professor Feigenbaum and collaborators. The handbook is broad in scope, covering all of the important ideas, techniques, and systems developed during 20 years of research in AI in a series of articles. Each is about four pages long and is a description written for non-AI specialists and students of AI. The handbook will be published in three volumes, the first of which is now on the market published by William Kaufmann, Inc. The AI Handbook effort is described in more detail in Section II.A.1.2 on page 99 and an outline of the current contents of the handbook can be found in Appendix B.

2) The AGE project which is an attempt to isolate inference, control, and representation techniques from previously developed knowledge-based programs; reprogram them for domain independence; write a rule-based interface that will help a user understand what the package offers and how to use the modules; and make the package available to other members of the AIM community. A more detailed description of progress on the AGE package can be found in Section II.A.1.1 on page 91.

It should be noted that SUMEX is providing only partial support for these projects with complementary support coming from an ARPA contract to the Heuristic Programming Project.

I.A.3.8 Resource Operations Statistics

The following data give an overview of various aspects of SUMEX-AIM resource usage. There are five sub-sections containing data respectively for:

1) Overall resource loading data

2) Relative system loading by community

3) Individual project and community usage

4) Network usage data

5) System reliability data

1. <u>Overall resource loading data</u>

The following plots display several different aspects of system loading over the life of the project. These include total CPU time delivered per month, the peak number of jobs logged in, and the peak load average. The monthly "peak" value of a given variable is the average of the daily peak values for that variable during the month. Thus, these "peak" values are representative of <u>average</u> monthly loading maxima and do not reflect the largest excursions seen on individual days, which are much higher.

These data show well the continued growth of SUMEX use and the self-limiting saturation effect of system load average, especially after installation of our overload controls early in 1978. Since late 1976, when the dual processor capacity became fully used, the peak daily load average has remained between about 5.5 and 6. This is a measure of the user capacity of our current hardware configuration and the mix of AI programs.

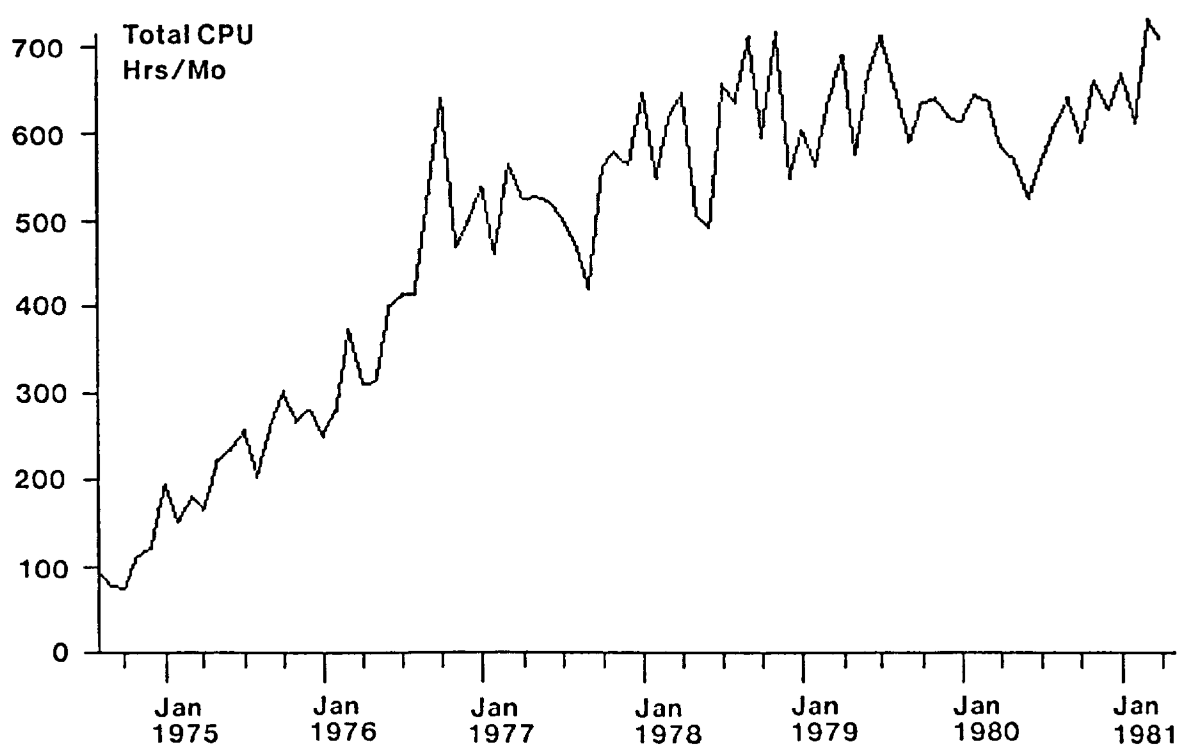

Figure 6. Total CPU Time Consumed by Month

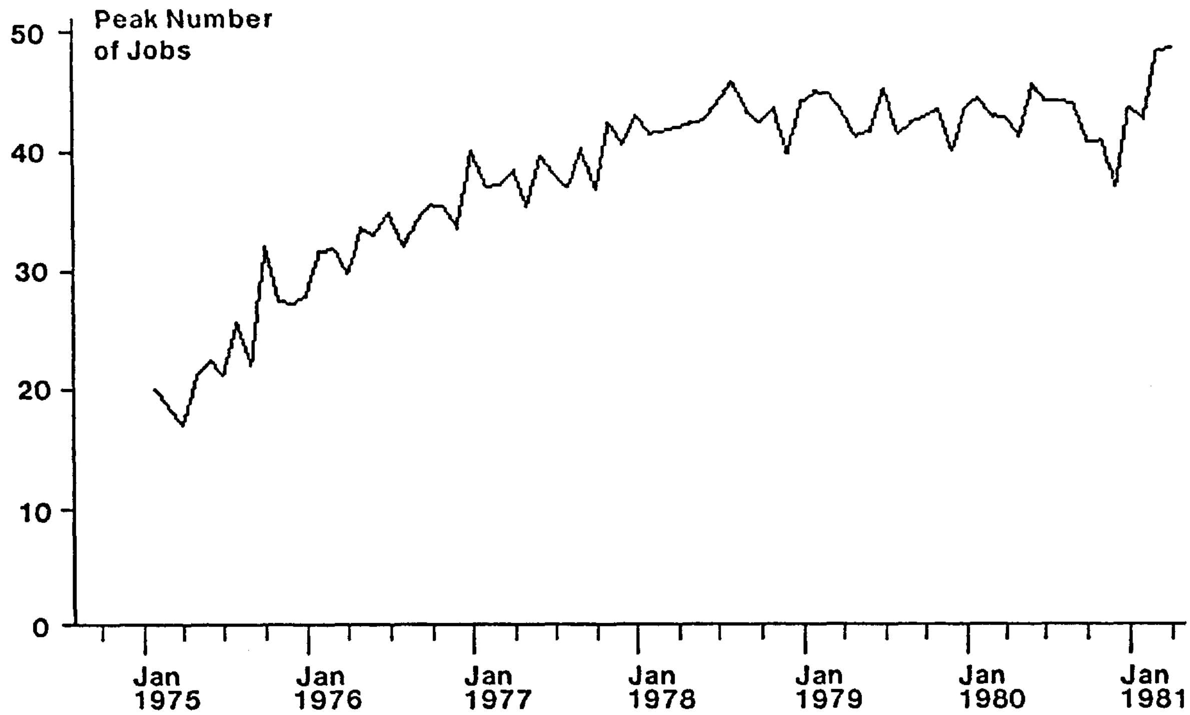

Figure 7. Peak Number of Jobs by Month

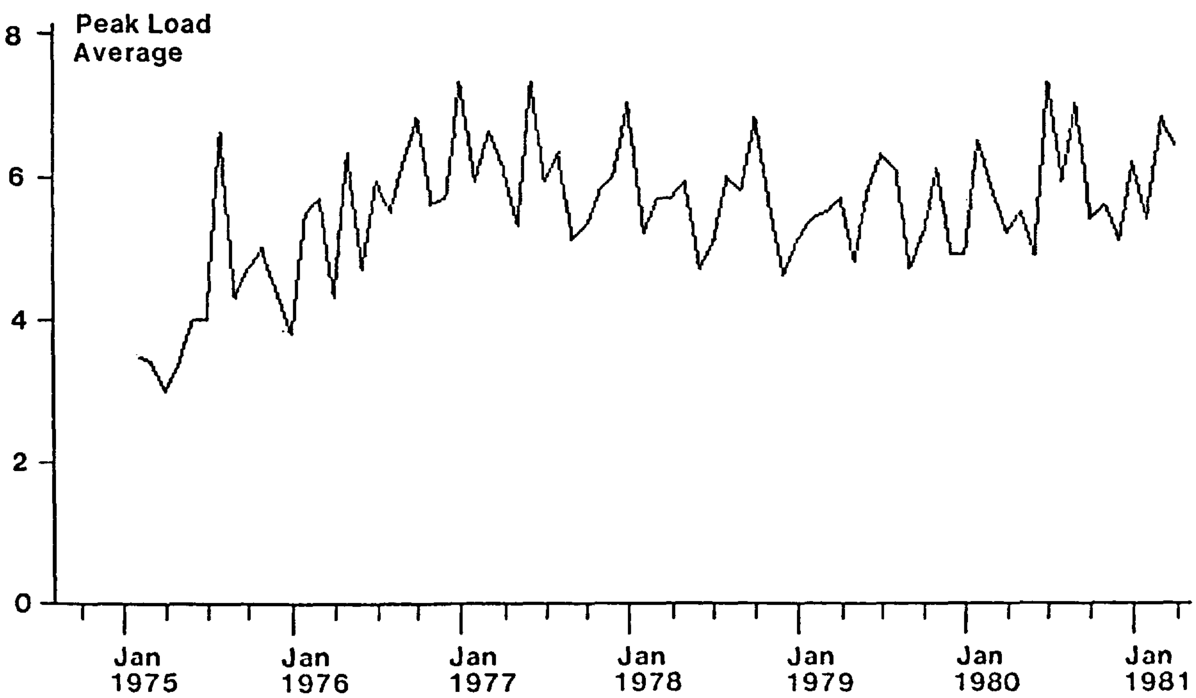

Figure 8. Peak Load Average by Month

2. Relative System Loading by Community

The SUMEX resource is divided, for administrative purposes, into 3 major communities: user projects based at the Stanford Medical School, user projects based outside of Stanford (national AIM projects), and common system development efforts. As defined in the resource management plan approved by BRP at the start of the project, the available system CPU capacity and file space resources are divided between these communities as follows:

Stanford	40%
AIM	40%
Staff	20%

The "available" resources to be divided up in this way are those remaining after various monitor and community-wide functions are accounted for. These include such things as job scheduling, overhead, network service, file space for subsystems, documentation, etc.

The monthly usage of CPU and file space resources for each of these three communities relative to their respective aliquots is shown in the plots in Figure 9 and Figure 10. Terminal connect time is shown in Figure 11. It is clear that the Stanford projects have held an edge in system usage despite our efforts at resource allocation and the substantial voluntary efforts by the Stanford community to utilize non-prime hours. This reflects the maturity of the Stanford group of projects relative to those getting started on the national side and has correspondingly accounted for much of the progress in AI program development to date.

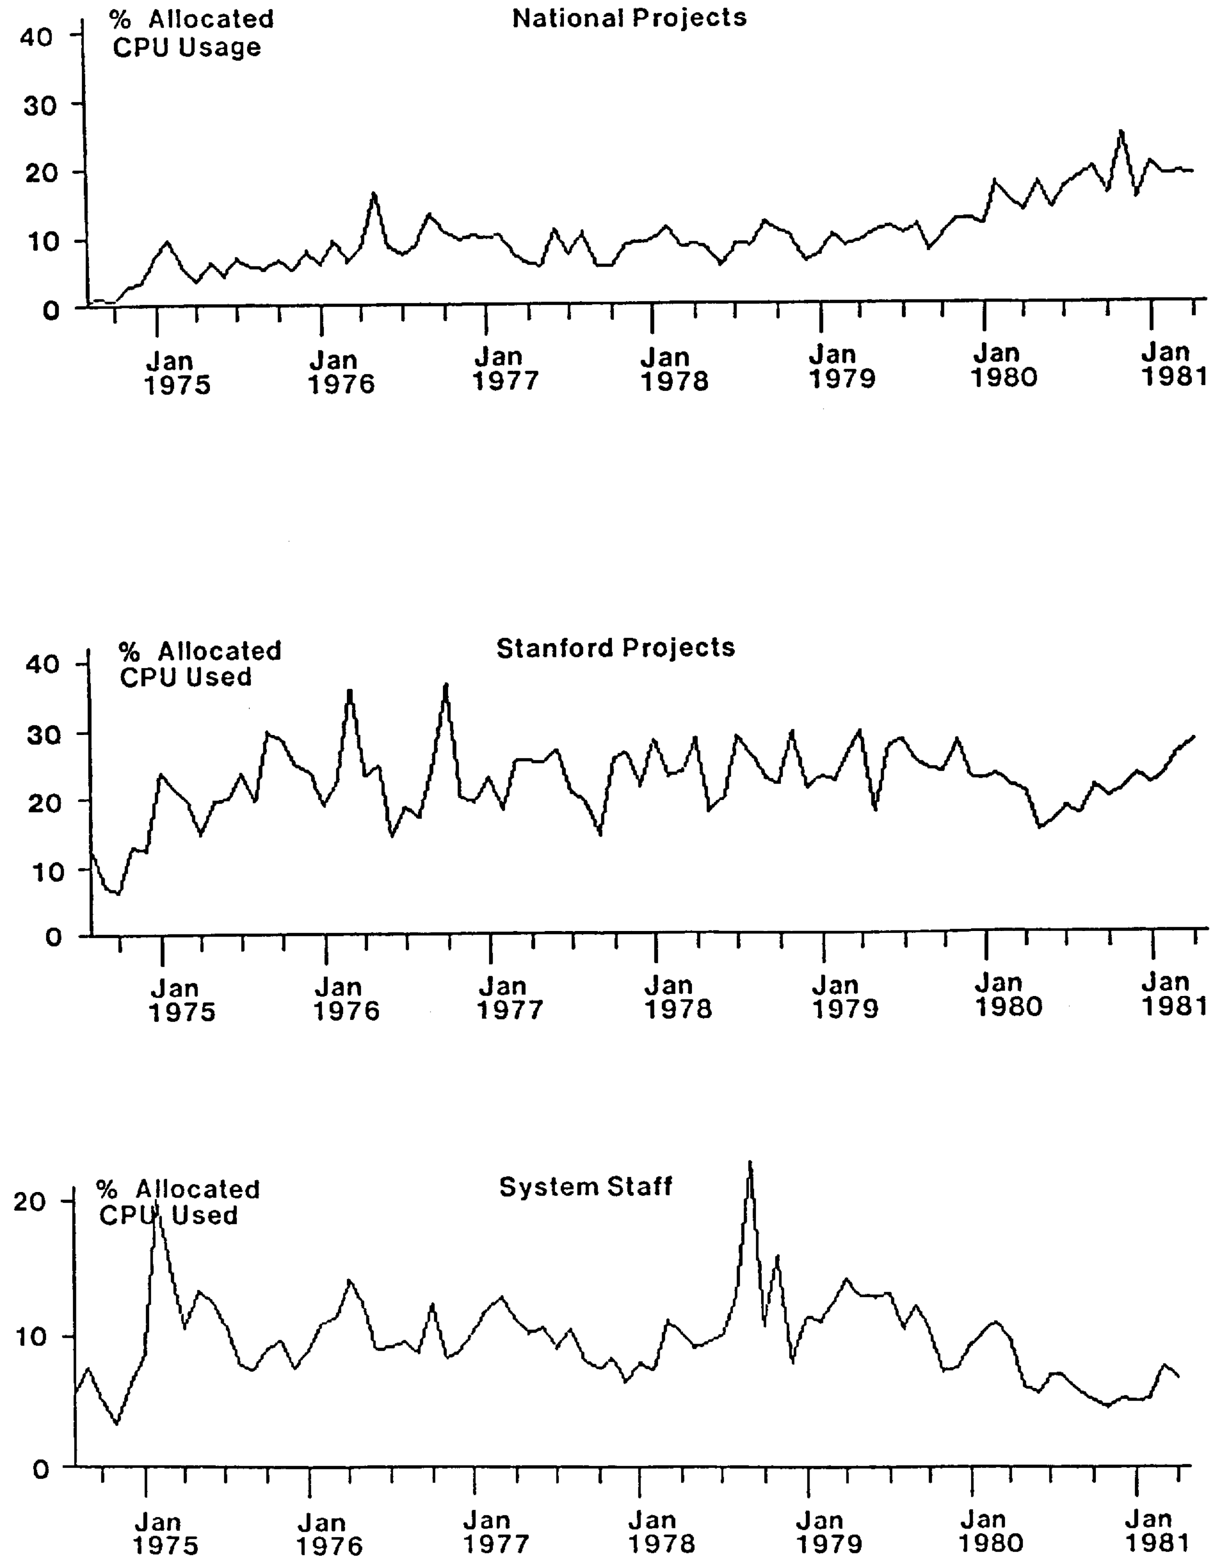

Figure 9. Monthly CPU Usage by Community

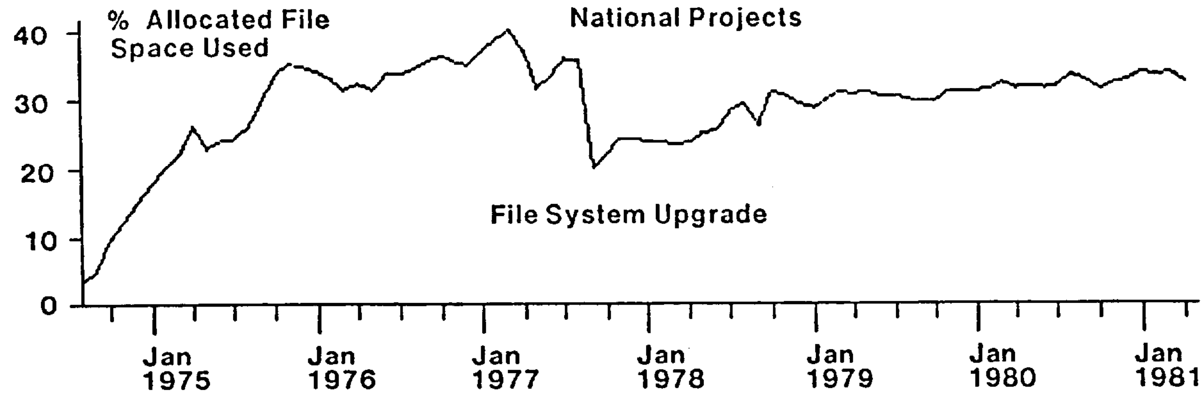

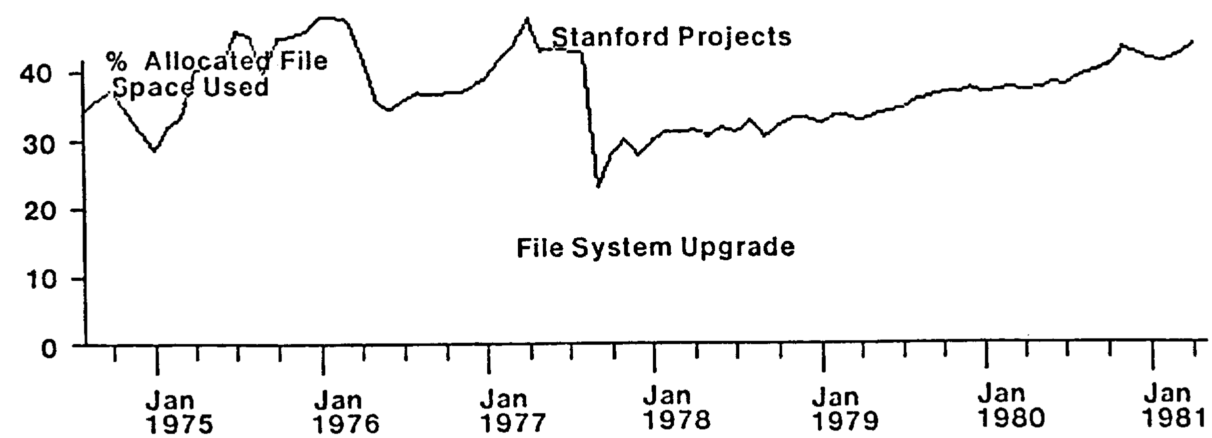

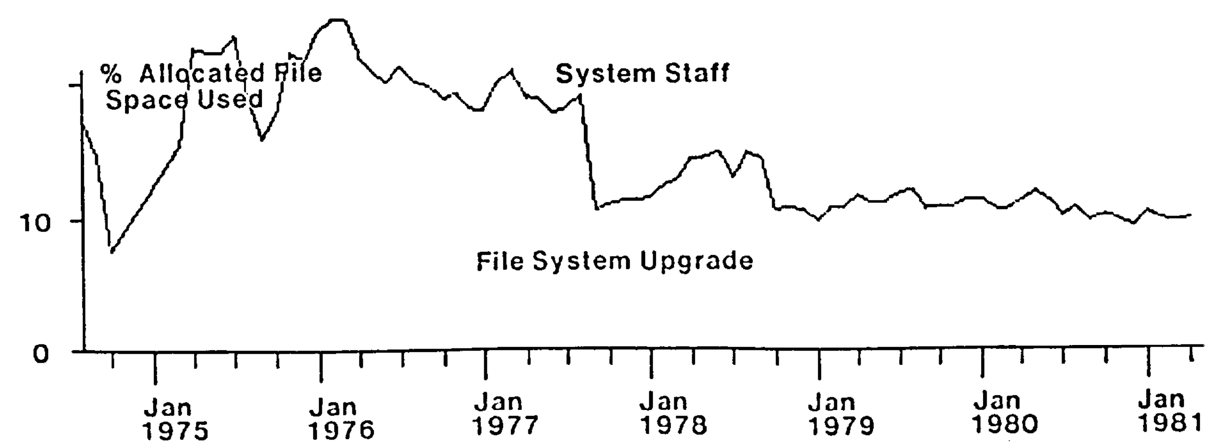

Figure 10. Monthly File Space Usage by Community

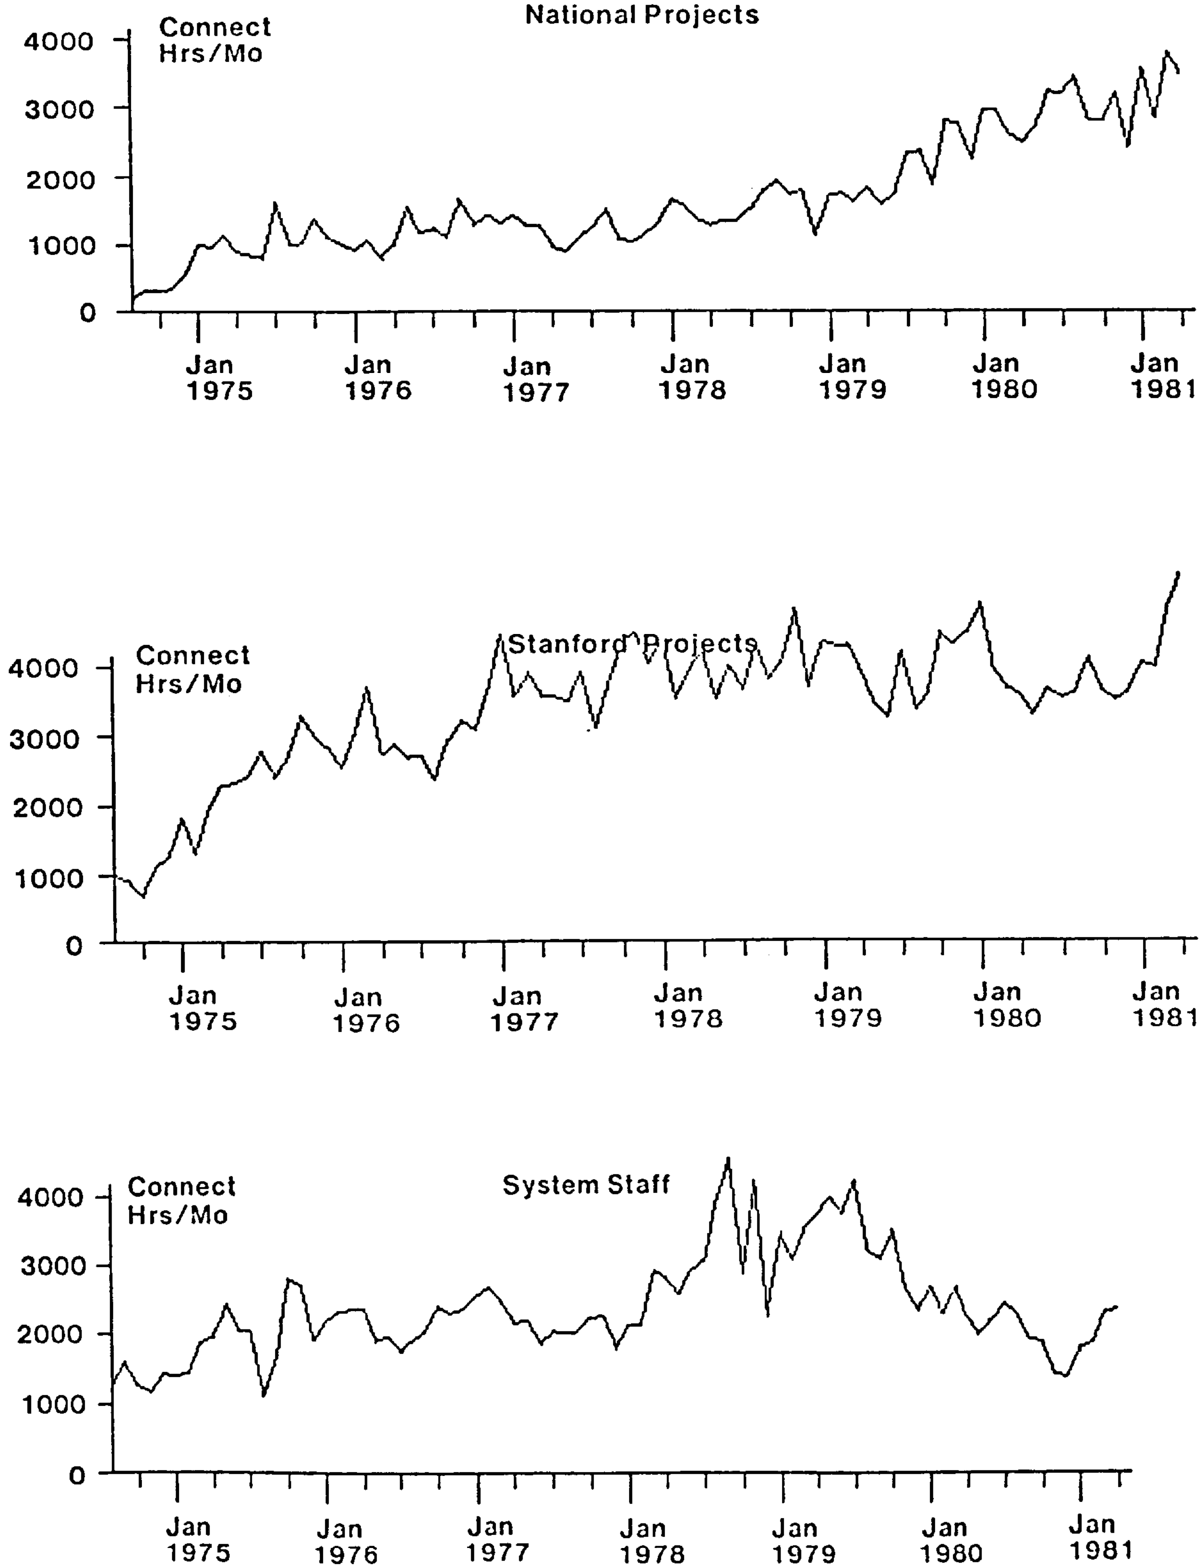

Figure 11. Monthly Terminal Connect Time by Community

3. Individual Project and Community Usage

The table following shows cumulative resource usage by project during the past grant year. The entries include a summary of the operational funding sources (outside of SUMEX-supplied computing resources) for currently active projects, total CPU consumption by project (Hours), total terminal connect time by project (Hours), and average file space in use by project (Pages, 1 page = 512 computer words). These data were accumulated for each project for the months between May 1980 and April 1981.

Again the well developed use of the SUMEX resource by the Stanford community can be seen. It should be noted that the Stanford projects have voluntarily shifted a substantial part of their development work to non-prime time hours which is not explicitly shown in these cumulative data. It should also be noted that a significant part of the DENDRAL, MYCIN, AGE, AI Handbook, and MOLGEN efforts, here charged to the Stanford aliquot, support development efforts dedicated to national community access to these systems. The actual demonstration and use of these programs by extramural users (e.g., the GENET community) is charged to the national community in the "AIM USERS" category, however.

Several of the projects admitted to the National AIM community use the Rutgers-AIM resource as their home base. We do not explicitly list these projects in this annual report covering the Stanford SUMEX-AIM resource. We do record information about the Rutgers resource itself, however, and note its separate resource status with the flag "[Rutgers-AIM]".

Resource Use by Individual Project - 5/80 through 4/81

National AIM Community	CPU (Hours)	Connect (Hours)	File Space (Pages)
1) ACT Project "Acquisition of Cognitive Procedures" John R. Anderson, Ph.D. Carnegie-Mellon Univ. ONR N00014-77-C-0242 9/78-9/80 $175,000 NSF IST-80-15357 2/81-2/84 $186,000	95.81	1214.90	2362.
2) SECS Project "Simulation & Evaluation of Chemical Synthesis" W. Todd Wipke, Ph.D. U. California, Santa Cruz NIH RR-01059-03S1 7/80-12/81 $36,949 NIH/NCI N01-CP-75816 1/80-7/81 $74,394	837.86	11968.57	10239
3) Mod Human Cogn Project "Hierarchical Models of Human Cognition" Walter Kintsch, Ph.D. Peter G. Polson, Ph.D. University of Colorado NIE-G-78-0172 9/80-8/81 $46,537 NIMH MH-15872-9-13 6/80-5/81 $32,880 ONR N00014-78-C-0433 6/80-5/81 $60,000 ONR N00014-78-C-0165 1/80-6/81 $85,000	150.08	2084.62	898
4) CADUCEUS Project "Clinical Decision Systems Research Resource" Jack D. Myers, M.D. Harry E. Pople, Jr., Ph.D. University of Pittsburgh NIH RR-01101-04 7/80-6/81 $465,199 NLM LM03710-01 7/80-6/81 $148,458 NLM LM03589-01 7/80-6/81 $32,750	344.92	5975.17	8365

5)	SOLVER Project "Problem Solving Expertise" Paul E. Johnson, Ph.D. William B. Thompson, Ph.D. University of Minnesota NSF SE079-13036 NICHD T36-HD-17151 NICHD HD-01136 NSF/BNS-77-22075 NLM/NSF proposals pending	.55	13.94	9
6)	PUFF-VM Project "Biomedical Knowledge Engineering in Clinical Medicine" John J. Osborn, M.D. Inst. Medical Sciences, San Francisco Edward A. Feigenbaum, Ph.D. Stanford University NIH GM-24669 9/78-8/81 $164,000 (*) Renewal pending	97.01	5651.45	3785
7)	SCP Project "Simulation of Cognitive Processes" James G. Greeno, Ph.D. Alan M. Lesgold, Ph.D. University of Pittsburgh NIE-G-80-0014 12/80-11/81 $2,627,067 ONR N00014-79-C-0215 10/80-9/81 $247,053 NSF/NIE SED78-22289 12/78-8/81 $149,967	5.13	236.97	931
8)	*** [Rutgers-AIM] *** Rutgers Project "Computers in Biomedicine" Saul Amarel, D.Sc. NIH RR-00643 12/80-11/81 $495,079	12.87	339.45	8653

9) AIM Pilot Projects			
AI-COAG	4.02	96.89	672
EXCHANGE	7.64	100.89	60
HEADMED	.66	29.03	770
KRL	.57	6.99	268
MDX	.53	30.96	28
MELANOMA	.99	20.11	9
MISL	.73	14.42	813
SPA	1.42	30.10	251
SPEECH	11.10	224.32	537
UI	8.51	379.92	46
	-----	-------	----
AIM Pilot Totals	36.17	933.63	3459
10) AIM Administration	8.98	449.06	4141
11) AIM Users on Stanford Projects			
AGE	2.44	51.90	27
DENDRAL	62.59	905.71	1056
HMF	35.06	1026.88	2735
HPP	5.30	118.23	115
MOLGEN	411.85	5757.19	1095
MYCIN	7.44	620.40	105
AIM-Associates	1.65	37.00	153
Guest (all projects)	12.88	135.08	239
	------	-------	----
AIM User Totals	569.21	8652.39	5529
	------	--------	-----
Community Totals	2158.59	37520.15	48371

Stanford Community	CPU (Hours)	Connect (Hours)	File Space (Pages)
1) AGE Project (Core) "Generalization of AI Tools" Edward A. Feigenbaum, Ph.D. Dept. Computer Science ARPA MDA-903-80-C-0107 (**) (partial support)	388.99	4625.17	3575
2) AI Handbook Project (Core) Edward A. Feigenbaum, Ph.D. Dept. Computer Science ARPA MDA-903-80-C-0107 (**) (partial support)	31.56	1680.34	2554
3) DENDRAL Project "Resource Related Research: Computers in Chemistry" Carl Djerassi, Ph.D. Dept. Chemistry NIH RR-00612-12 5/81-4/82 $237,387	554.47	8170.34	16366
4) EXPEX Project "Expert Explanation" Edward H. Shortliffe, M.D., Ph.D. Depts. Medicine/Computer Science ONR NR 049-479 1/81-12/81 $140,825	26.99	1163.37	513
5) MOLGEN Project "Experiment Planning System for Molecular Genetics" Edward A. Feigenbaum, Ph.D. Bruce G. Buchanan, Ph.D. Laurence H. Kedes, M.D. Douglas L. Brutlag, Ph.D. Depts. Computer Science/ Medicine/Biochemistry NSF ECS-8016247 10/80-9/81 $146,582 (*)	325.62	6146.69	6734

6) MYCIN Projects "Computer-based Consult. in Clin. Therapeutics" Bruce G. Buchanan, Ph.D. Edward H. Shortliffe, M.D., Ph.D. Depts. Medicine/Computer Science NSF MCS-79-03753 7/79-3/81 $146,152 ONR/ARPA N00014-79-C-0302 3/79-3/82 $396,325 Kaiser Fdn. 7/79-12/80 $20,000 NLM LM-03395 7/80-6/81 $47,845 NLM LM-00048 7/80-6/81 $39,107	706.62	11373.94	13261
7) Protein Struct Modeling "Heuristic Comp. Applied to Prot. Crystallog." Edward A. Feigenbaum, Ph.D. Dept. Computer Science NSF MCS-79-33666 12/79-11/81 $35,318	89.26	1194.98	3916
8) RX Project Depts. Computer Science/Medicine Robert L. Blum, M.D. Gio C.M. Wiederhold, Ph.D. NLM New Invest. 7/79-6/82 $90,000 NCHSR 4/79-3/81 $35,000	64.15	1683.05	2222
9) Stanford Pilot Projects			
DECIDER (E. Johnson)	.27	3.70	0
STRUCT (Abarbanel)	12.41	317.34	87
SCANR (Brinkley)	36.81	758.69	771
	------	-------	----
Stanford Pilot Totals	49.49	1079.73	858
10) Stanford and HPP Assoc.	196.49	9538.63	11609
	-------	--------	-----
Community Totals	2433.64	46656.24	61608

SUMEX Staff	CPU (Hours)	Connect (Hours)	File Space (Pages)
1) Staff	602.53	20648.45	9571
2) System Associates	78.42	3081.22	4205
3) Misc. Usage	.16	1.84	770
	-------	---------	------
Community Totals	681.11	23731.51	14546

System Operations	CPU (Hours)	Connect (Hours)	File Space (Pages)
1) Operations	2230.32	94975.93	69094
	=======	=========	======
Resource Totals	7503.66	202883.83	193619

(*) Award includes indirect costs.

(**) Supported by a larger ARPA contract MDA-903-80-C-0107 awarded to the Stanford Computer Science Department:

	Current Year (10/80-11/15/81)	Total Award (10/79-9/82)
Heuristic Programming Project	$ 538,262	$1,613,588
VLSI/CAD Network-based Graphics Development Resource	214,851	685,374
	-----------	----------
Total award	$ 753,113(*)	$2,298,962(*)

4. Network Usage Statistics

The plots in Figure 12 and Figure 13 show the monthly network terminal connect time for TYMNET and ARPANET. This forms the major billing component for SUMEX-AIM TYMNET usage. The terminal connect time does not reflect the time spent in file transfers and mail forwarding.

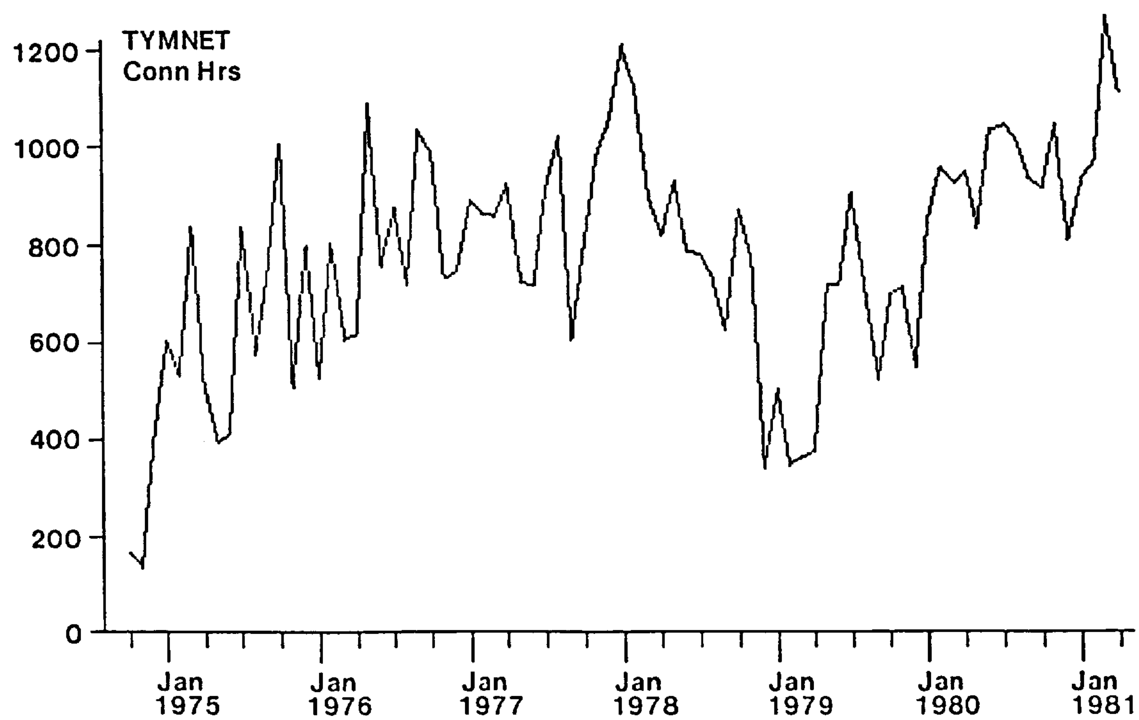

Figure 12. TYMNET Terminal Connect Time

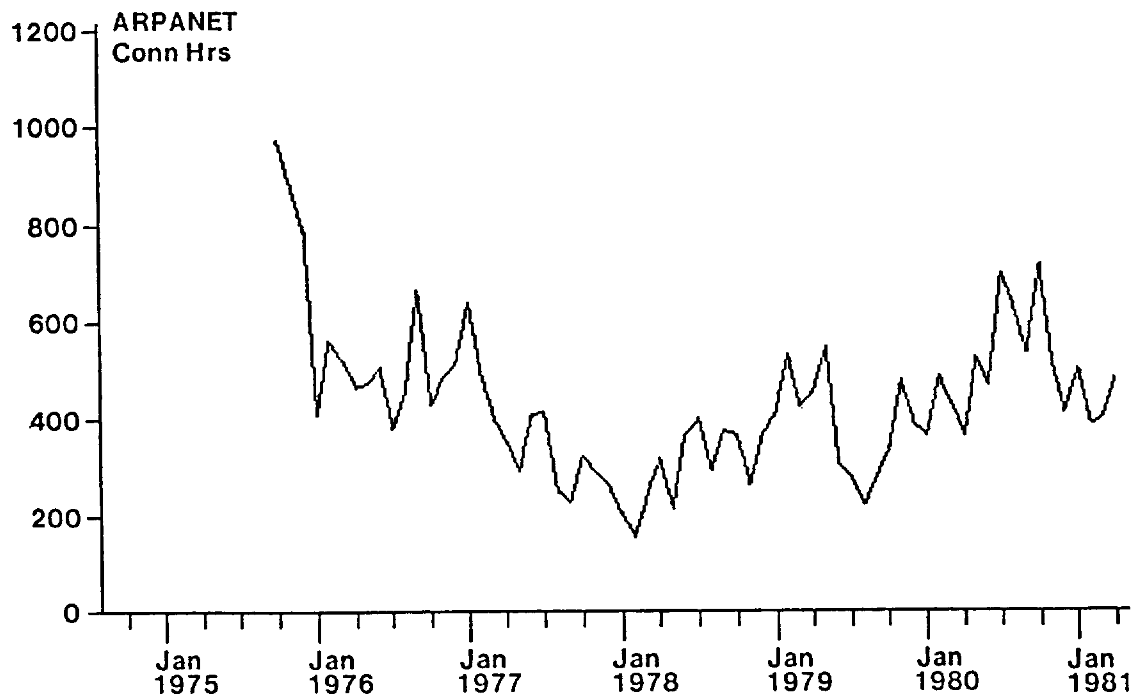

Figure 13. ARPANET Terminal Connect Time

5. System Reliability

System reliability has been very good on average with several periods of particular hardware or software problems. The table below shows monthly system reloads and downtime for the past year. It should be noted that the number of system reloads is greater than the actual number of system crashes since two or more reloads may have to be done within minutes of each other after a crash to repair file damage or to diagnose the cause of failure.

				1980						1981		
	MAY	JUN	JUL	AUG	SEP	OCT	NOV	DEC	JAN	FEB	MAR	APR
RELOADS												
Hardware	3	2	2	8	8	16	2	0	6	2	4	0
Software	0	3	2	4	1	2	1	1	3	2	1	4
Environmental	0	1	0	3	1	0	0	1	0	0	0	1
Operator Error	0	0	1	0	0	0	1	1	0	0	2	1
Unknown Cause	2	0	0	0	0	0	0	0	0	0	0	0
Totals	5	6	5	15	10	18	4	3	9	4	7	6
DOWNTIME (Hrs)												
Unscheduled	2	17	6	18	7	11	1	2	6	2	4	6
Scheduled	30	14	79	17	17	28	17	7	15	4	4	14
Totals (Hrs)	32	31	95	35	24	39	18	9	21	6	8	20

TABLE 1. System Reliability by Month

I.A.3.9 SUMEX Staff Publications

The following are publications for the SUMEX staff and include papers describing the SUMEX-AIM resource and on-going research as well as documentation of system and program developments. Many of the publications documenting SUMEX-AIM community research are from the individual collaborating projects and are detailed in their respective reports (see Section II on page 89). Publications for the AGE and AI Handbook core research projects are given there.

[1] Carhart, R.E., Johnson, S.M., Smith, D.H., Buchanan, B.G., Dromey, R.G., and Lederberg, J, Networking and a Collaborative Research Community: A Case Study Using the DENDRAL Programs, ACS Symposium Series, Number 19, Computer Networking and Chemistry, Peter Lykos (Editor), 1975.

[2] Levinthal, E.C., Carhart, R.E., Johnson, S.M., and Lederberg, J., When Computers Talk to Computers, Industrial Research, November 1975

[3] Wilcox, C. R., MAINSAIL - A Machine-Independent Programming System, Proceedings of the DEC Users Society, Vol. 2, No. 4, Spring 1976.

[4] Wilcox, Clark R., The MAINSAIL Project: Developing Tools for Software Portability, Proceedings, Computer Application in Medical Care, October, 1977, pp. 76-83.

[5] Lederberg, J. L., Digital Communications and the Conduct of Science: The New Literacy, Proc. IEEE, Vol. 66, No. 11, Nov 1978.

[6] Wilcox, C. R., Jirak, G. A., and Dageforde, M. L., MAINSAIL - Language Manual, Stanford University Computer Science Report STAN-CS-80-791 (1980).

[7] Wilcox, C. R., Jirak, G. A., and Dageforde, M. L., MAINSAIL - Implementation Overview, Stanford University Computer Science Report STAN-CS-80-792 (1980).

Mr. Clark Wilcox also chaired the session on "Languages for Portability" at the DECUS DECsystem10 Spring '76 Symposium.

In addition, a substantial continuing effort has gone into developing, upgrading, and extending documentation about the SUMEX-AIM resource, the SUMEX-TENEX system, and the many subsystems available to users. These efforts include a number of major documents (such as SOS, PUB, TENEX-SAIL, and MAINSAIL manuals) as well as a much larger number of document upgrades, user information and introductory notes, an ARPANET Resource Handbook entry, and policy guidelines.

I.A.3.10 Future Plans

Our plans for the next grant year are based on those approved by the council review of our recent five-year renewal application scheduled to begin in August 1980. In addition to the specific plans for next grant year (discussed in some earlier sections too), we present a summary below of our overall objectives for the next five-year period to serve as a foundation for future reports. Near and long term objectives and plans for individual collaborating projects are discussed in Section II beginning on page 89.

The goals of the SUMEX-AIM resource are long term in supporting basic research in artificial intelligence, applying these techniques to a broad range of biomedical problems, experimenting with communication technologies to promote scientific interchange, and developing better tools and facilities to carry on this research.

Just as the tone of our renewal proposal derives from the continuing long-term research objectives of the SUMEX-AIM community, our approach derives from the methods and philosophy already established for the resource. We will continue to develop useful knowledge-based software tools for biomedical research based on innovative, yet accessible computing technologies.

For us it is important to make systems that work and are exportable. Hence, our approach is to integrate available state-of-the art hardware technology as a basis for the underlying software research and development necessary to support the AI work.

SUMEX-AIM will retain its broad community orientation in choosing and implementing its resources. We will draw upon the expertise of on-going research efforts where possible and build on these where extensions or innovations are necessary. This orientation has proved to be an effective way to build the current facility and community.

We have built ties to a broad computer science community; have brought the results of their work to the AIM users; and have exported results of our own work. This broader community is particularly active in developing technological tools in the form of new machine architectures, language support, and interactive modalities.

Toward a More Distributed Resource

The initial model for SUMEX as a centralized resource was based on the high cost of powerful computing facilities and not being able to duplicate them readily. This role is evolving with the introduction of more compact and inexpensive computing technology. Our future goals are guided by community needs for more computing capacity and improved tools to build more effective expert systems and to test operational versions of AI programs in real-world settings. In order to meet these needs, we must take advantage of a range of newly developing machine architectures and systems. As a result, SUMEX-AIM will become more a distributed community resource with heterogeneous computing facilities tethered to each other through communications media. Many of these machines will be located physically near the projects or biomedical scientists using them.

We have actively supported proposals from the more mature AIM projects for additional computing facilities tailored to their particular needs and designed to free the main SUMEX resource for new, developing applications projects. To date, the Rutgers resource has acquired a DEC 2050 facility, part of which is allocated for AIM usage; the "Simulation of Cognitive Processes" project has acquired a VAX which supports their needs; and the "Caduceus" (INTERNIST) project is acquiring a VAX to support experimental clinical testing of their program. Our future plans anticipate an even broader diversification of computing resources to meet the need of the AIM community.

The Continuing Role of SUMEX-Central

Even with more distributed computing resources, the central resource will continue to play an important role as a communication crossroads, as a research group devoted to integrating the new software and hardware technologies to meet the needs of medical AI applications, as a spawning ground for new application projects, and as a base for local AI projects. A key challenge will be to maintain the scientific community ties that grew naturally out of the previous colocation within a central facility.

Summary of Five-year Objectives

The following outlines the specific objectives of the SUMEX-AIM resource during the follow-on five year period. Note that these objectives cover only the resource nucleus; near and long-term objectives for individual collaborating projects are discussed in their respective reports in Section II. Specific aims are broken into three categories; 1) resource operations, 2) training and education, and 3) core research.

Resource Operations

1) Maintain the vitality of the AIM community. We will continue to encourage and explore new applications of AI to biomedical research and improve mechanisms for inter- and intra-group collaborations and communications. While AI is our defining theme, we may entertain exceptional applications justified by some other unique feature of SUMEX-AIM essential for important biomedical research. To minimize administrative barriers to the community-oriented goals of SUMEX-AIM and to direct our resources toward purely scientific goals, we plan to retain the current user funding arrangements for projects working on SUMEX facilities. User projects will fund their own manpower and local needs; will actively contribute their special expertise to the SUMEX-AIM community; and will receive an allocation of computing resources under the control of the AIM management committees. There will be no "fee for service" charges for community members. We will also continue to exploit community expertise and sharing in software development; and to facilitate more effective information sharing among projects.

2) Continue to provide effective computational support for AIM community goals. Our efforts will be to extend the support for artificial intelligence research and new applications work; to develop new computational tools to support more mature projects; and to facilitate testing and research dissemination of nearly operational programs. We will continue to operate and develop the existing KI-10/2020 facility as the nucleus of the resource. We will acquire additional equipment to meet developing community needs for more capacity, larger program address spaces, and improved interactive facilities. New computing hardware technologies becoming available now and in the next few years will play a key role in these developments and we expect to take the lead in this community for adapting these new tools to biomedical AI needs. We plan the phased purchase of two VAX computers to provide increased computing capacity and to support large address space LISP development, a 2000M byte file server to meet file storage needs, and a number of single-user "professional workstations" to experiment with improved human interfaces and AI program dissemination.

3) Provide effective and geographically accessible communication facilities to the SUMEX-AIM community for effective remote collaborations, communications among distributed computing nodes,

and experimental testing of AI programs. We will retain the current ARPANET and TYMNET connections for at least the near term and will actively explore other advantageous connections to new communications networks and to dedicated links.

Training and Education

Our goals during the follow-on period for assisting new and established users of the SUMEX-AIM resource are a continuation of those adopted for the previous grant term. Collaborating projects are responsible for the development and dissemination of their own AI programs. The SUMEX resource will provide community-wide support and will work to make resource goals and AI programs known and available to appropriate medical scientists. Specific aims include:

1) Provide documentation and assistance to interface users to resource facilities and programs. We will continue to exploit particular areas of expertise within the community for developing pilot efforts in new application areas.

2) Continue to allocate "collaborative linkage" funds to qualifying new and pilot projects to provide for communications and terminal support pending formal approval and funding of their projects. These funds are allocated in cooperation with the AIM Executive Committee reviews of prospective user projects.

3) Continue to support workshop activities including collaboration with the Rutgers Computers in Biomedicine resource on the AIM community workshop and with individual projects for more specialized workshops covering specific application areas or program dissemination.